arXiv:1403.7326v1 [math.LO] 28 Mar 2014

IMAGINARIES AND DEFINABLE TYPES IN ALGEBRAICALLY CLOSED VALUED FIELDS

EHUD HRUSHOVSKI

This manuscript is largely an exposition of material from [1], [2] and [3], regarding definable types in the model completion of the theory of valued fields, and the classification of imaginary sorts. The proof of the latter is based here on definable types rather than invariant types, and on the notion of *generic reparametrization*; it allows a more conceptual view than we had when [1] was written. I also try to bring out the relation to the geometry of [3] - stably dominated definable types as the model theoretic incarnation of a Berkovich point.

The text is based on notes from a class entitled *Model Theory of Berkovich Spaces*, given at the Hebrew University in the fall term of 2009, and retains the flavor of class notes. Thanks to Adina Cohen, Itai Kaplan, and Daniel Lowengrub for comments. Most recently, Will Johnson went through the notes with great care; he is due thanks for numerous textual improvements as well as some highly perceptive mathematical comments and corrections. He further discovered a considerable simplification of the proof of elimination of imaginaries, based on definable types and their coding in $\mathcal{O}$-submodules of finite dimensional K-spaces, but shortcutting the decomposition theorem of definable types, Theorem 5.3; this proof, I hope, will appear separately.

The material was discussed in my talk in the Valuation Theory meeting in El Escorial in 2011. The slides for this talk can be found in [4].

0.1. **Notation.** We will use a universal domain for a given theory, usually the theory ACVF defined below. This is a highly saturated and highly homogeneous model, denoted $\mathbb{U}$. Small subsets of $\mathbb{U}$ are denoted by A, B, ... Definable subsets of $\mathbb{U}$ are denoted by X, Y, ..., and sometimes D. If M is a model containing the parameters used to define X, $X(M)$ denotes the interpretation of X in M. If A is a substructure of a model and $x_1, \ldots, x_n$ are tuples from the model, then $A(x_1, \ldots, x_n)$ denotes the definable closure of $A, x_1, \ldots, x_n$.

When working with valued fields, the valued field itself is denoted K, the residue field is denoted k, the valuation ring is denoted $\mathcal{O}$, the maximal ideal is denoted $\mathcal{M}$, and the value group is denoted Γ. The residue map is $\mathrm{res} : \mathcal{O} \to \mathrm{k}$, and the value map is $\mathrm{val} : K \to \Gamma \cup \{\infty\}$. The value group is written additively, so that $\mathcal{O} = \{x \in K : v(x) \geq 0\}$. ACVF is the theory of non-trivially valued algebraically closed valued fields.

Let B_n denote the group of invertible upper triangular matrices. The group of elements of B_n with entries in a given ring R is denoted $B_n(R)$. We will also write B_n for $B_n(K)$. U_n is the group of matrices in B_n with 1's on the diagonal. D_n is the group of diagonal matrices, so that $B_n = D_n U_n$.

1. Definable types

1.1. **Definable types.** Let L_Y be the set of formulas of T in variables from Y, up to T-equivalence. A *definable type* $p(x)$ is a family of Boolean retraction $L_{x,Y}$ to L_Y (for any finite set of variables Y), compatible with inclusions $Y \subset Y'$. It is denoted: $\phi \mapsto (d_p x)\phi$. Thus $(d_p x)\phi$ is a formula with (at most) the same y-variables but without the free variable x; it is analogous to quantifiers, but simpler; one says: for generic $x \models p$, ϕ holds.

Given a definable type p and a substructure A of $M \models T$, we let

$$p|A = \{\phi(x,a) : a \in A^l, M \models (d_p x)\phi(a)\}$$

So we can think of a definable type as a compatible family of types, given systematically over all base sets.

1.2. **Examples, notation.** While the development is at first abstract, we will give examples from ACVF, the theory of algebraically closed valued fields. K denotes the field, $\mathcal{O}$ the valuation ring, Γ the value group, val the valuation map, res the residue map into the residue field k.

1.3. **Pushforward of definable types.** Let $f : X \to Y$ be an A- definable function, and p an A- definable type on X. Define $q = f_* p$ by:

$$(d_q y)\theta(y,u) = (d_p x)\theta(f(x),u)$$

Excercise. For any B containing A we have: $(f_* p)|B = tp(f(c)/B)$ where $c \models p|B$.

1.4. **Product of definable types.** If p and q are two A-definable types, then the product $p(x) \otimes q(y)$ is defined by

$$(d_{p\otimes q}(x,y))\theta(x,y,u) = (d_q y)(d_p x)\theta(x,y,u).$$

If B contains A, then $(c_1,c_2) \models p \otimes q|B$ if and only if $c_2 \models q|B$ and $c_1 \models p|B(c_2)$.

1.5. **Orthogonality.** A definable type $q(x)$ is *constant* if $(d_q x)(x = y)$ has a solution.

Excercise. In this case, $(d_q x)(x = y)$ has a unique solution a; and a is the unique realization of $q|B$, for any B over which q is defined.

Definition 1.6. p is *orthogonal to* Γ if for any $\mathbb{U}$-definable function f into Γ, $f_* p$ is constant.

Equivalently, by considering coordinate projections, any $\mathbb{U}$-definable function f into Γ^n is constant. We will use this definition for the value group, which eliminates imaginaries; otherwise we would instead consider definable functions f into Γ^{eq}.

1.7. **Stable embeddedness.** A sort D is *stably embedded* if any $\mathbb{U}$-definable subset of D^m is $D(\mathbb{U})$-definable.

In ACVF, both Γ and k are *stably embedded*; this is an immediate consequence of quantifier-elimination in the standard three-sorted language (See Theorem 2.1.1 (iii) in [1], or the first paragraph of the Appendix.) It suffices to consider atomic formulas, with some variables from Γ and some from other sorts. Any atomic formula $\phi(x_1, \ldots, x_n, y_1, \ldots, y_n)$ with x_i in Γ, $y_j \in VF$, has the form: $\theta(x_1, \ldots, x_n, \mathrm{val}(h_\nu(y)))$. So $\phi(x, b)$ defines the same set as $\theta(x, d)$ where $d = \mathrm{val}h(b)$. Similarly for k and res$h$, with h a rational function.

Orthogonality of p to Γ can also be stated as follows: Let $B' = B(\gamma)$ be generated over B by some realizations of Γ. Then $p|B$ implies $p|B'$.

1.8. **Domination.**

Lemma 1.9. *Let $f : X \to Y$ be an A-definable function. Let q be an A-definable type on Y, and let p_A be a type over A on X. Assume: for any $B \geq A$ there exists a unique type p_B such that p_B contains p_A, and $f_* p_B = q|B$. Then there exists a unique A-definable type p such that for all B, $p|B = p_B$.*

Proof. More generally, let us say a type $p_{\mathbb{U}}$ over $\mathbb{U}$ is *definably generated* over A if it is generated by a partial type of the form $\cup_{(\phi,\theta)\in S} P(\phi, \theta)$, where S is a (small) set of pairs of formulas $(\phi(x, y), \theta(y))$ over A, and $P(\phi, \theta) = \{\phi(x, b) : \theta(b)\}$.

It sufices to show that if $p_{\mathbb{U}}$ is definably generated over A, then $p_{\mathbb{U}}$ is definable over A, i.e. $\{b : \phi(x, b) \in p_{\mathbb{U}}\}$ is A-definable for each A-formula $\phi(x, y)$.

Let $\phi(x, y)$ be any formula. From the fact that $p_{\mathbb{U}}$ is definably generated it follows easily that $\{b : \phi(x, b) \in p_{\mathbb{U}}\}$ is an $\bigvee$-definable set over A, i.e. a union of A-definable sets. Indeed, $\phi(x, b) \in p_{\mathbb{U}}$ if and only if for some $(\phi_1, \theta_1), \ldots, (\phi_m, \theta_m) \in S$, $(\exists c_1, \cdots, c_m)(\theta_i(c_i) \wedge (\forall x)(\bigwedge_i \phi_i(x, c) \implies \phi(x, b))$. Applying this to $\neg\phi$, we see that the complement of $\{b : \phi(x, b) \in p_{\mathbb{U}}\}$ is also $\bigvee$-definable. Hence $\{b : \phi(x, b) \in p_{\mathbb{U}}\}$ is A- definable.

□

Definition 1.10. In the situation of the lemma, p is said to be *dominated by q* via f

In the situation of the lemma, p is said to be *dominated by q* via f. More precisely:

Definition 1.11. p is dominated by q via f if there is some A over which p, q, and f are defined, such that for every $B \geq A$, $(q|B)(f(x)) \cup (p|A)(x) \vdash (p|B)(x)$.

In general, when p, q, f are A-definable, one can visualize that p is dominated by q over some bigger set B, but not over A. When A is a model, this does not happen, nor will it occur in our setting of stable domination (see Remark 3.10). (Thanks to Will Johnson for this remark.)

Example 1.12. (ACVF) Let $X = \mathcal{O}, Y = \mathrm{k}, f = \mathrm{res}$. Let q be the generic type of k, i.e. $q|B$ is generated by: $y \in \mathrm{k}, y \notin V$ for any finite B-definable V. Then $x \in \mathcal{O}, f(x) \models q|B$ generates a complete type $p|B$ over B. This is called the *generic type of* $\mathcal{O}$.

Exercise 1.13. Show that $p|B$ is complete. For any polynomial $\sum b_i x^i$ over B, show that $\mathrm{val}(\sum b_i x^i) = \min_i \mathrm{val}(b_i)$ for x realizing $p|B$. In particular, p is orthogonal to Γ.

Example 1.14. Let $\mathcal{M} = \{x : \mathrm{val}(x) > 0\}$ be the maximal ideal. Let $f(x) = \mathrm{val}(x)$. Let $q(x)$ be the type just above 0 in Γ. Then q dominates via f a definable type $p_{\mathcal{M}}$, the generic type of $\mathcal{M}$.

Example 1.15. ($\mathrm{ACVF}_{0,0}$). Let $a_0, a_1, \ldots \in \mathbb{Q}$. Let $\mathrm{val}(t) > 0$. Let $p_0(x, y)$ consist of all formulas (over $\mathbb{Q}(t)$)

$$\mathrm{val}(y - \sum_{k=0}^{n} a_k (xt)^k \geq n\mathrm{val}(t))$$

Then $p_0(x, y) + (p_{\mathcal{O}}|\mathbb{U})(x)$ generates a complete type $p|\mathbb{U}$, provided $\sum a_n x^n$ is transcendental.

Let $X = \mathcal{O} \times \mathcal{O}$, $Y = \mathrm{k}$, $f(x, y) = \mathrm{res}(x)$. Then p is dominated by the generic type of k, via f.

To prove the domination, say $\mathrm{val}(t) = 1$. First let M be a valued field extension of $\mathbb{Q}(t)^{alg}$ such that $\mathbb{Z}$ is cofinal in $\mathrm{val}(M)$. We prove domination over M.

Generalizing the construction, allow $a_n \in \mathcal{O}_M$, $a = \sum a_k X^k$, and define p_0^a to consist of all formulas:

$$\mathrm{val}(y - \sum_{k=0}^{n} a_k (xt)^k) \geq n$$

For a fixed, write $p_0 = p_0^a$.

Let $c \models p_{\mathcal{O}}|M$. First suppose $p_0(c, 0)$ holds. Then $\min_{i \leq n} \mathrm{val}(a_i) + i = \mathrm{val}(\sum_{k=0}^{n} a_k (ct)^k) \geq n$. So $\mathrm{val}(a_i) \geq n - i$. Letting $n \to \infty$ (and using $a_i \in M$) we see that $a_i = 0$; so $a = 0$.

Next suppose just that $p_0(c, d)$ holds for some $d \in M(c)^{alg}$. So $F(c, d) = 0$ for some polynomial $F \in \mathcal{O}_M[x, y]$. Let $a' = F(x, a(x))$ be the power series obtained by substituting $a(x)$ for y. Let $p_0' = p_0^{a'}$. Then $p_0'(c, 0)$ holds. Hence by the previous paragraph, $a' = 0$, so a is algebraic.

Otherwise, $p_0(c, y)$ defines an infinite intersection b of balls over $M(c)$, with no algebraic point. Hence b contains no nonempty $M(c)$-definable subset ($M(c)^{alg} \models$

$ACVF$, so any nonempty $M(c)$-definable set does have an algebraic point.) So $p_0 + tp(c/M)$ generates a complete type over $M(c)$, as promised.

We can take M to be maximally complete; this suffices to show that $p|M$ is stably dominated.

Now if N is a valued field extension of M with $\mathrm{res}(N) = \mathrm{res}(M)$, then $p|M \vdash p|N$, hence $p_0(x, y) + p_0|M$ already generates $p|N$.

But any valued field extension of $\mathbb{Q}(t)^{alg}$ can be obtained in this way (taking such an M, N and then a subextension.) This proves the domination statement in the example.

1.16. **Density of definable types.** We consider the following extension property for a definable set D over a base set A, possibly including imaginaries.

(E(A,D)): Either $D = \emptyset$, or there exists a definable type p on D (over $\mathbb{U}$) such that p has a finite orbit under $Aut(\mathbb{U}/A)$.

Say T has property E if $E(A, D)$ holds for all A, D. In Lemma 5.2 below, we will see that ACVF has property (E).

We say that a substructure B of $\mathbb{U}$ is a *canonical base* for an object p constructed from $\mathbb{U}$ if for any $\sigma \in Aut(\mathbb{U})$, $\sigma(p) = p$ iff $\sigma(b) = b$ for all $b \in B$.

Lemma 1.17. *Let T be a theory with property (E), and assume any definable type (in the basic sorts) has a canonical base in certain imaginary sorts $S_1, S_2, \ldots$. Then T admits elimination of imaginaries to the level of finite subsets of products of the S_i.*

1.18. **Definable types on Γ^n.** Let Γ be a divisible ordered Abelian group. Recall that the theory of divisible ordered Abelian groups has quantifier-elimination (a result whose roots go back to Fourier.)

We will consider projections $\phi^a : \Gamma^n \to \Gamma$, $\phi^a(x) = a \cdot x$, where $a \in \mathbb{Q}^n \smallsetminus (0)$.

We say two definable types p, q are *orthogonal* if there is a set A over which p and q are defined, such that for any $B \geq A$, $p(x)|B \cup q(y)|B$ generates a complete type in the variables x, y.

A definable type p in Γ^n has a *limit* if there is some $c \in \Gamma^n$ such that for every $\mathbb{U}$-definable open neighborhood U of c, the formula $x \in U$ is in $p|\mathbb{U}$.

Lemma 1.19. *Let p be a definable type of Γ^n, over A. Then up to a change of coordinates by a rational $n \times n$ matrix, p decomposes as the join of two orthogonal definable types p_f, p_i, such that p_f has a limit in Γ^m, and $\phi^a_* p_i$ has limit point $\pm\infty$ for any $a \in \mathbb{Q}^n \smallsetminus (0)$.*

Proof. Let $\alpha_1, \ldots, \alpha_k$ be a maximal set of linearly independent vectors in $\mathbb{Q}^n$ such that the image of p under $(x_1, \ldots, x_n) \mapsto \alpha_i \cdot x$ has a limit point in Γ [1] Let $\beta_1, \ldots, \beta_l$ be a maximal set of vectors in $\mathbb{Q}^n$ such that for any/every model M and for $x \models p|M$, $\alpha_1 x, \ldots, \alpha_k x, \beta_1 x, \cdots, \beta_l x$ are linearly independent over

[1] Equivalently, the image of p under $x \mapsto (\alpha_1 x, \ldots, \alpha_k x)$ has a limit point in Γ^k.

$M(\alpha_1 x, \ldots, \alpha_k x)$ If $a \models p|M$, let $a' = (\alpha_1 a, \ldots, \alpha_k a)$, $a'' = (\beta_1 a, \ldots, \beta_k a)$. For $\alpha \in \mathbb{Q}(\alpha_1, \ldots, \alpha_k)$, the element αa is bounded between elements of M. On the other hand each βa ($\beta \in \mathbb{Q}(\beta_1, \ldots, \beta_k) \smallsetminus \{0\}$) satisfies $\beta a > M$ or $\beta a < M$. For if $m \leq \beta a'' \leq m'$ for some $m \in M$, since $\mathrm{tp}(\beta a''/M)$ is definable it must have a finite limit, contradicting the maximality of k. It follows that $\mathrm{tp}(\alpha a/M) \cup \mathrm{tp}(\beta a/M)$ extends to a complete 2-type, namely $\mathrm{tp}((\alpha a, \beta a)/M)$; in particular $\mathrm{tp}(\alpha a + \beta a/M)$ is determined; from this, by quantifier elimination, $\mathrm{tp}(a'/M) \cup \mathrm{tp}(a''/M)$ extends to a unique type in $k + l$ variables. So $\mathrm{tp}(a'/M)$, $\mathrm{tp}(a''/M)$ are orthogonal. After some sign changes in a'', so that each coordinate is $> M$, the lemma follows. □

Lemma 1.20. (1) *Let p, p' be definable types on Γ^n. If $\phi^a_* p = \phi^a_* p'$ for each a, then $p = p'$.*

(2) *Let p be a definable type on Γ^n. If $\phi^a_* p$ is 0-definable for each a, then p is 0-definable.*

Proof. (1) Any formula $\phi(x, y)$ is a Boolean combination of formulas $a \cdot x + b \cdot y > c$ (or $= c$). The definition of such a formula is determined by $\phi^a_* p$.

(2) Let σ be an automorphism, $p' = \sigma(p)$; we have to show that $p' = p$. This follows from (1). □

Lemma 1.21. *Let p be a definable type of Γ^n. For $c \in \Gamma^n$, let $\alpha^c(x) = x + c$. Then for some $c \in \Gamma^n$, $\alpha^c_* p$ is 0-definable.*

Proof. A linear change of coordinates (with $\mathbb{Q}$-coefficients) does not effect this statement. So we may assume the conclusion of Lemma 1.19 holds. Translating the p_f part by $-\lim p_f$, we may assume p_f has limit $0 \in \Gamma^m$. It follows that for any $a \in \mathbb{Q}^n \smallsetminus (0)$, $\phi^a_* p$ has limit 0 or $\pm\infty$. There are only five definable 1-types with this property, all 0-definable. Hence by Lemma 1.20(2), p is 0-definable. □

2. Algebraic lemmas on valued fields

The material in this section is classical, going back in part to Ostrowsky and Kaplansky; see the book by F.V.-Kuhlmann `http://math.usask.ca/ fvk/Fvkbook.htm`.

Definition 2.1. An extension $L \leq L'$ of valued fields is *immediate* if L, L' have the same value group and residue field.

K is *maximally complete* if it has no proper immediate extensions.

Exercise 2.2. Let K be an algebraically closed valued field, L a valued field extension, $t \in L$. Assume $L = K(t)$ as a field. Since any element of $K[t]$ is a product of linear factors, the valuation on L is determined by $v(t - a)$ for $a \in K$. Then one of the following holds:

- $v(t - a) = \gamma \notin \Gamma(K)$ for some $a \in K$. Show that $\Gamma(L) = \Gamma(K)(\gamma)$, $\mathrm{k}(L) = \mathrm{k}(K)$.

- $v(t-a) \in \Gamma(K)$ for all $a \in K$, and $v(t-a)$ takes a maximal value $v(b)$ at some $a \in K$. Show that $\mathrm{k}(L) = \mathrm{k}(K)(e)$ where $e = \mathrm{res}((t-a)/b)$.
- $v(t-a) \in \Gamma(K)$ for all $a \in K$, and a maximum is not attained. Show that $K(t)$ is an immediate extension.

Lemma 2.3. *Let L/K be an extension of valued fields. Then $tr.deg._{\mathrm{res}K}\mathrm{res}L + \dim_{\mathbb{Q}}(\mathrm{val}(L)/\mathrm{val}(K)) \leq tr.deg._K(L)$.*

Proof. This reduces to the case that L/K is generated by one element. In this case L/K is algebraic or $L = K(t)$ is a rational function field. In the algebraic case, res$L$ is a finite extension of res$K$ (of some degree e) and $\mathrm{val}(L)/\mathrm{val}(F)$ is finite (of some order f; in fact we have $ef \leq [L:K]$.) In case $L = K(t)$, we may assume K is algebraically closed, since passing to this case will not lower the left hand side; and Ex. 2.2 applies. □

Lemma 2.4. *Let K denote a valued field, with algebraically closed residue field* k *and divisible value group A. Assume K is maximally complete,*

- *K is algebraically closed.*
- *K is* spherically complete*, i.e. any set of balls, linearly ordered by inclusion, has nonempty intersection.*

Proof. (1) This follows from Lemma 2.3: algebraic extensions are immediate since the value group and residue field have no proper finite extensions. (2) Let b_i be a set of balls, indexed by a linear ordering I. If $\cap_i b_i = \emptyset$, then for any $a \in K$ we have $a \notin b_i$ for large i, and it follows that $\alpha(a) = v(a-c)$ is constant for $c \in b_i$. Define a valuation on $K(t)$ by $v(t-a) = \alpha(a)$. Then by Ex. 2.2 this is an immediate extension, a contradiction. □

Any valued field K has a maximally complete immediate extension, of cardinality at most $2^{|K|}$.

2.5. **Valued vector spaces.** A *valued* vector space over valued field K is a triple $(V, \Gamma(V), v)$, with V a K-space, $\Gamma(V)$ a linearly ordered set $\Gamma(V)$ with an action $+ : \Gamma(K) \times \Gamma(V) \to \Gamma(V)$, order-preserving in each variable, and v a map $v : V \smallsetminus (0) \to \Gamma(V)$ with $v(a+b) \geq \min(v(a), v(b))$ and $v(cb) = v(c) + v(b)$ for $a, b \in V, c \in K$.

If $a_1, \ldots, a_n$ are elements of V with $v(a_1), \ldots, v(a_n)$ in distinct $\Gamma(K)$-orbits, it follows that $a_1, \ldots, a_n$ are linearly independent over K. In particular if V is finite-dimensional, $\Gamma(V)$ can only consist of finitely many $\Gamma(K)$-orbits.

By a *ball* in V we mean a set of the form $\{b \in V : v(a-b) \geq \alpha\}$. V is *spherically complete* if any set of balls, linearly ordered by inclusion, has nonempty intersection.

A set $a_1, \ldots, a_n$ of elements of V is called *separated* if for all $c_1, \ldots, c_n \in K$, we have

$$v\left(\sum c_i a_i\right) = \min_i (v(c_i) + v(a_i))$$

Such a set is in particular linearly independent.

If $V = K^n$ is a valued K-space with a separated basis, a ball for V is just a product of balls of K, so V is spherically complete if K is.

If V is a valued K-space with a spherically complete subspace $W \leq V$, and $a \in V$, then the set $\{v(w-a) : w \in W\}$ attains a maximum, because for each $\gamma \in \Gamma(V)$, the set $\{w \in W : v(w-a) \geq \gamma\}$ is either empty or a ball in W.

Lemma 2.6. *Let K be a spherically complete valued field, V a finite-dimensional K-space. Then V has a separated basis.*

Proof. Let $a_1, \ldots, a_m$ be a maximal separated set, U the subspace generated by $a_1, \ldots, a_m$. Then U has a separated basis, so it is spherically complete. If $U = V$ we are done. Otherwise, let $a \in V \setminus U$. Consider the possible values $v(u-a)$, $u \in U$. Since U it is spherically complete, so there must be a maximal value among these. Replacing a by $a - u$ with $v(a-u)$ maximal, we may assume $v(a) \geq v(a-u)$ for all $u \in U$. In this case, $a_1, \ldots, a_m, a$ is separated. For given $c_1, \ldots, c_m$, we have $v(\sum c_i a_i) = \min_i v(c_i a_i) = \gamma$ say. It suffices to see that $v(\sum c_i a_i + a) \leq \min(\gamma, v(a))$; this follows from the strong triangle inequality when $\gamma \neq v(a)$, and from $v(a) \geq v(\sum c_i a_i + a)$ when $\gamma = v(a)$. □

2.7. **Induced** k**-spaces.** Let V be a valued K-space, and $\alpha \in \Gamma(V)$. Then $\Lambda_\alpha = \{a \in V : v(a) \geq \alpha\}$ is an $\mathcal{O}$-submodule, and $\Lambda^o_\alpha = a \in V : v(a) > \alpha\}$ is an $\mathcal{O}$-submodule containing $\mathcal{M}\Lambda_\alpha$. Let $V_\alpha = \Lambda_\alpha / \Lambda^o_\alpha$; this is a $\mathrm{k} = \mathcal{O}/\mathcal{M}$-space, finite-dimensional if V is.

Let $h : U \to V$ be a homomorphism of valued K-spaces; meaning there is also a map $h : \Gamma(U) \to \Gamma(V)$ of $\Gamma(K)$-sets, with $h(\alpha) < h(\beta)$ when $\alpha < \beta$, and $v(h(a)) = h(v(a))$. Then h induces a homomorphism $U_\alpha \to V_{h(\alpha)}$ for each α.

2.8. **Tensor products.** Let U, V be valued K-spaces. Consider K-spaces $(W, \Gamma(W))$ and maps

$$h : U \otimes V \to W, \quad + : \Gamma(U) \times \Gamma(V) \to \Gamma(W)$$

such that $v(h(a \otimes b)) = v(a) + v(b)$ and $\Gamma(U) \times \Gamma(V) \to \Gamma(W)$ is order-preserving in each variable.

Then for each $\alpha \in \Gamma(U), \beta \in \Gamma(V)$ we have an induced homomorphism

$$U_\alpha \otimes V_\beta \to W_{\alpha+\beta}$$

Lemma 2.9. *Let K be spherically complete, and let U, V be valued K-spaces. Let E be a divisible ordered Abelian group with $\Gamma(K)$-action, and assume $\Gamma(U), \Gamma(V) \leq E$ and $E \cong \Gamma(U) \times_{\Gamma(K)} \Gamma(V)$, i.e. if $\alpha + \beta = \alpha' + \beta'$ then for some $\gamma \in \Gamma(K)$, $\gamma + \alpha = \alpha'$ and $\gamma + \beta' = \beta$. Then there exists a unique $(W, h : W \to E)$ (up to a unique isomorphism) such that:*

1. *For any $(\alpha, \beta) \in \Gamma(U) \times \Gamma(V)$, the induced homomorphism $U_\alpha \otimes V_\beta \to W_{\alpha+\beta}$ is injective.*

Proof. To prove uniqueness we have to show that h is injective, and determine $v(h(x))$ for all $x \in U \otimes V$. Write $x = \sum_{i=1}^{n} a_i \otimes b_i$ where $(a_1, \ldots, a_n)$ are separated. Then it suffices to show:

Claim . $h(x) \neq 0$, and $v(h(x)) = \min_i v(a_i) + v(b_i)$.

By grouping the terms according to the value of $v(a_i) + v(b_i)$, it suffices to prove the claim when $v(a_i) + v(b_i)$ is constant. In this case by assumption there exists $\gamma_i = v(c_i), c_i \in K$ with $v(a_i) = v(a_1) + \gamma, v(b_i) = v(b_1) - \gamma$. Replacing a_i by a_i/c_i and b_i by $b_i c_i$, we may assume $v(a_i) = v(a_1) = \alpha, v(b_i) = v(b_i) = \beta$. So $a_i \in \Lambda_\alpha^U, b_i \in \Lambda_\beta^V$. Since $a_1, \ldots, a_n$ are separated, the images $\bar{a}_i$ of the a_i in U_α are linearly independent. The images $\bar{b}_i$ of the b_i in V_β are nonzero. Hence $\sum \bar{a}_i \otimes \bar{b}_i \neq 0 \in U_\alpha \otimes V_\beta$. Since h induces an injective map into $W_{\alpha+\beta}$ it follows that $v(h(\sum_i a_i \otimes b_i)) = \alpha + \beta$.

With uniqueness proved, functoriality is clear and so it suffices to prove existence in the finite dimensional case. This is easily done by choosing a separated basis and following the recipe implicit above. □

Proposition 2.10. *Let K be a spherically complete valued field, L_1, L_2 valued field extensions, within a valued field extension N generated by $L_1 \cup L_2$. Assume $\Gamma(K) = \Gamma(L_1)$, and $\mathrm{k}(L_1)$ is linearly disjoint from $\mathrm{k}(L_2)$ over $\mathrm{k}(K)$. Then the structure of the valued field N is uniquely determined given L_1 and L_2.*

Proof. It suffices to show that the natural map $h : L_1 \otimes L_2 \to N$ is injective and that $v(h(x))$ is determined for $x \in L_1 \otimes L_2$, since passage to the field of fractions is clear using $v(x/y) = v(x) - v(y)$. Let W be the image of h. Then we are in the setting of Lemma 2.9, (1) holds, and (2) is clear since $\Gamma(U) = \Gamma(K)$. For the same reason, (3) reduces to the case $\alpha = 0$. Suppose $\mathrm{val}(b_1) = \min_i \mathrm{val}(b_i)$, without loss of generality. We have $h(\sum_i a_i \otimes b_i) = b_1 h(\sum_i a_i \otimes (b_i/b_1))$, so we may take $\beta = 0$ too. In this case (3) amounts to the linear disjointess assumption. The corollary now follows from the lemma. □

Proposition 2.10 will imply that any definable type orthogonal to Γ is dominated by its images in k. We did not use Lemma 2.9 in full generality; using it we could deduce that any definable type is dominated by its images in Γ and in k. We will in fact require a stronger statement, of stable domination relative to Γ. The algebraic content consists of the lemma below.

Let L_1, L_2 be two valued field extensions L_1, L_2 of a valued field K, contained in a valued field extension N of K, and such that $L_1 \cup L_2$ generates N. As in Lemma 2.10, we will say that the interaction between L_1, L_2 is uniquely determined (given some conditions) if whenever N' is another valued field extensions of K, and $j_i : L_i \to N'$ are valued K-algebra homomorphisms (satisfying the same conditions), then there exists a (unique) valued K-algebra embedding $j : N \to N'$ with $j|L_i = j_i$.

It is easy to see that condition (2) below does not depend on the choice of Z.

Proposition 2.11. *Let K be a spherically complete valued field, L_1, L_2 valued field extensions, within a valued field extension N generated by $L_1 \cup L_2$. Let $\mathrm{k}_0 = \mathrm{k}(K)$, $\mathrm{k}_i = \mathrm{k}(L_i)$, $\mathrm{k}_{12} = \mathrm{k}_1\mathrm{k}_2$. Then the interaction of L_1, L_2 is uniquely determined assuming the following conditions.*

1. $\Gamma(L_1) \subseteq \Gamma(L_2)$.
2. *Let Z be a $\mathbb{Q}$-basis for $\Gamma(L_1)/\Gamma(K)$; for $z \in Z$ let $a_z \in L_1$ and $b_z \in L_2$ have $v(a_z) = v(b_z) = z$, and let $c_z = a_z/b_z$. Assume the elements $\mathrm{res}(c_z)$ form an algebraically independent set over k_{12}.*
3. *$\mathrm{k}_1, \mathrm{k}_2$ are linearly disjoint over k_0.*

Proof. As in Lemma 2.10, it suffices to show that the natural map $h : L_1 \otimes L_2 \to N$ is injective and that $v(h(x))$ is determined for $x \in L_1 \otimes L_2$. Write $x = \sum_{i=1}^{n} a_i \otimes b_i$, with (a_i) separated. We claim that $v(x) = \min_i v(a_i) + v(b_i) \in \Gamma(L_2)$. As before we may assume $v(a_i) + v(b_i) = \gamma$ does not depend on i. Moreover since $\gamma = v(c)$ for some $c \in L_2$, dividing b_i by c we may assume $\gamma = 0$. The subgroup of $\Gamma(L_2)$ generated by the $v(a_i)$ is finitely generated; let $d_1, \ldots, d_l$ be a minimal set of generators of this group modulo $\Gamma(K)$. Let a'_z, b'_z, c_z be as in condition (2), so that $v(a'_z) = v(b'_z) = d_z$ for $z = 1, \ldots, l$, $c_z = a'_z/b'_z$, and the elements $\mathrm{res}(c_z)$ are algebraically independent over k_{12}. For each i, there exists $m = (m_1, \ldots, m_l) \in \mathbb{Z}^l$ with $\sum m_z d_z = v(a_i)$. Write $m(i)$ for this m, and $(a')^m$ for $\Pi_z (a'_z)^{m_z}$, and similarly for b' and c. Let $A_i = a_i/(a')^{m(i)}$, $B_i = b_i (b')^{m(i)}$. Then

$$\sum a_i b_i = \sum A_i B_i c^{-m(i)}$$

We have to show that this has valuation zero, i.e. that

$$\sum_i \mathrm{res}(A_i)\mathrm{res}(B_i)\mathrm{res}(c)^{-m(i)} \neq 0$$

Since the $\mathrm{res}(c_z)$ are algebraically independent (2), it suffices to show that for a fixed value of $m \in \mathbb{Z}^l$, we have: $\sum_{m(i)=m} \mathrm{res}(A_i)\mathrm{res}(B_i) \neq 0$. But this follows from (3) as in Lemma 2.9. □

3. Stably dominated types

.

Definition 3.1. An A-definable type p is stably dominated if for some $B \geq A$, p is dominated over B by a definable map f into V for some finite-dimensional k-space V.

When the base A consists of elements of the valued field and Γ, it can be shown that f can be chosen to be A-definable. The space V is isomorphic to k^m over some larger B, but not necessarily over A. For instance, given $\alpha \in \Gamma$, let $\mathcal{O}\alpha = \mathcal{O}c$ where $\mathrm{val}(c) = \alpha$. Then $\mathcal{O}\alpha$ is a free $\mathcal{O}$-module, and $\mathcal{O}a/\mathcal{M}a$ is a one-dimensional k-space V_α.

Exercise 3.2. The generic type of the ball $\mathrm{val}(x) \geq \alpha$ is dominated by the map $r_\alpha : \mathcal{O}a \to V_\alpha$. However every α-definable map on $\mathcal{O}_\alpha$ into k is constant, if α is not a root of the valuation of some element of the prime field.

This special family of definable types will be the main object we will look at. For any definable set V, we will define $\widehat{V}$ to be the set of stably dominated types on V. Later, a topology will be defined on $\widehat{V}$; V will be dense in $\widehat{V}$, called the *stable completion* of V.

Theorem 3.3. *In ACVF, the following conditions on a definable type are equivalent:*

(1) *p is stably dominated.*
(2) *For all definable q, $p(x) \otimes q(y) = q(y) \otimes p(x)$*
(3) *p is symmetric: $p(x) \otimes p(y) = p(y) \otimes p(x)$.*
(4) *p is orthogonal to Γ.*

Proof. (1) implies (2): by domination it suffices to prove that $p(x) \otimes q(y) = q(y) \otimes p(x)$ for p on k^n. By stable embeddedness one reduces to the case that q too is on k^n.

(2) implies (3) is trivial.

(3) implies (4): Let f be a definable function into Γ. Then $q = f_* p$ is symmetric. But by considering the $q(u)$-definition of $u < v$ one sees that q must be constant.

(4) implies (1). Let M be a maximally complete valued field, with p definable over M. Let $a \models p|M$, $N = M(a)$. Then $\Gamma(N) = \Gamma(M)$ by orthogonality. By Proposition 3.5, a unique M-definable type extends $p|M$, and this type is stably dominated; this type must be p. □

Exercise 3.4. Let k be an algebraically closed field, V a finite-dimensional vector space over k, definable in some theory over a base A. We assume that the definable subsets of k^m are the constructible subsets. Let p_A be a type of elements of V, over A. Then there exists at most one A-definable type p such that $p|A = p_A$.

(Proof: p is the generic type of a unique Zariski-closed subset W of V; W must be A-definable; we must have $W \in p_A$ but no smaller subvariety is in p_A; this characterizes W and hence p.)

Proposition 3.5. *Let M be a maximally complete algebraically closed valued field, $N = M(a)$ a valued field extension. Let $\gamma = (\gamma_1, \ldots, \gamma_n)$ be a basis for $\Gamma(N)/\Gamma(M)$. Then there exists a unique $M(\gamma)$-definable type extending $tp(a/M(\gamma))$. This type is stably dominated.*

Proof. We have $\gamma_i = \mathrm{val}(c_i)$ for some $c_i \in N$. Let $e_i = r_{\gamma_i}(c_i) \in V_{\alpha_i}$ (see notation above.) Let $\alpha_1, \ldots, \alpha_m$ be a transcendence basis for $\mathrm{k}(N)$ over $\mathrm{k}(M)$. We have $\epsilon_i = g_i(a), \alpha_j = h_j(a)$ for some $M(\gamma)$-definable functions g_i, h_j. Let $V = \Pi_i V_{\alpha_i} \times \mathrm{k}^m$, $g = (g_1, \ldots, h_m)$. Let $q = q(v_1, \ldots, v_n, t_1, \ldots, t_m)$ be the generic type of

the k-space V; equivalently, letting $V_{i+n} = \mathrm{k}$, $q = q_1 \otimes \ldots \otimes q_{n+m}$ where q_i is the unique non-constant definable type on the 1-dimensional k-vector space V_i. Note that for any structure $B \geq M(\gamma)$, for $i \leq n$, if $e_i' \in V_i(B), e_i' \neq 0$ and $e_i \models q_i|B$ then e_i/e_i' is a well-defined element, realizing the generic type of k over B; hence if $(e_1, \ldots, e_n, \alpha_1, \ldots, \alpha_m) \models q|B$, then $(e_1/e_1', \ldots, e_n/e_n', \alpha_1, \ldots, \alpha_m)$ are algebraically independent over B. Note also in this situation that if if $\gamma_i = \mathrm{val}(b_i)$ with $b_i \in B$, then $e_i/e_i' = \mathrm{res}(c_i/b_i)$. By Proposition 2.11, there exists a unique type p_B extending $tp(a/M(\gamma))$ and with $g_* p_B = q|B$. By Lemma 1.9 there exists a unique $M(\gamma)$-definable type p with $p|B = p_B$ for all B.

By definition, p is dominated by g and hence stably dominated. If p is another $M(\gamma)$-definable type extending $tp(a/M(\gamma))$, let $q' = g_* p'$. Then q' is an $M(\gamma)$-definable type extending $tp(e_1, \ldots, \alpha_m)/M(\gamma)$. By Excercise 3.4 we have $q' = q$, and hence by the domination, $p' = p$. This proves the uniqueness of p. □

Discussion 3.6. Let V be an M-definable set, with $a \in V$. We will see below that $\widehat{V}$ can be viewed as a *pro-definable set*; i.e. an inverse limit of definable sets. In more detail: we will describe certain definable sets $\widehat{V}_d$ for $d \in \mathbb{N}$, and definable maps $\widehat{V}_{d+1} \to \widehat{V}_d$. (These maps can be taken to be surjective, but we will not use this fact here. The $\widehat{V}_d$ will be subsets of $K^m \times S_n$ for appropriate m, n, where S_n is the sort of lattices in K^n, described below.) Let $\varprojlim_{d \in \mathbb{N}} \widehat{V}_d$ be the set of sequences $c = (c_d : d \in \mathbb{N})$ such that $c_{d+1} \mapsto c_d$. Say $c \in \mathrm{dcl}(A)$ iff each $c_d \in \mathrm{dcl}(A)$. A definable map $f : X \to \varprojlim_{d \in \mathbb{N}} \widehat{V}_d$ means: a compatible system of definable maps $f_d : \Gamma^n \to \widehat{V}_d$.

For each $c \in \varprojlim_{d \in \mathbb{N}} \widehat{V}_d$ we will describe (canonically) a stably dominated type p_c. We will show that any stably dominated type on V equals p_c for a unique $c \in \widehat{V}$. It follows that $c \in \mathrm{dcl}(A)$ iff p_c is A-definable. We define $\widehat{V} = \varprojlim_{d \in \mathbb{N}} \widehat{V}_d$.

In this language, Proposition 3.5 states that there exists a pro-definable partial map $f : \Gamma^n \to \widehat{V}$ (over M) and $\gamma \in \Gamma^n$ such that with $c = f(\gamma)$, we have $\gamma \in M(a)$ and $a \models p_c|M(\gamma)$.

Thus $tp(a/M)$ can be understood in terms of (i) $tp(\gamma/M)$ and (ii) an M-definable function $\Gamma^n \to \widehat{V}$.

Exercise 3.7. Let r be an A-definable type, and let f be an A-pro-definable function into $\widehat{V}$, with $\mathrm{dom}(f) \in r|A$. For any B with $A \leq B$, let $a \models r|B$, $p = p_{f(a)}$, $c \models p|B(a)$. Show that $p_B = tp(c/B)$ does not depend on the choices, and that there exists a unique A-definable type p with $p|B = p_B$. We will refer to this type as $\int_r f$.

In particular, Proposition 3.5 and the discussion below it yield:

Exercise 3.8. Any M-definable type on V has the form $\int_r f$ for some M-definable type r on Γ^n, and some M-definable partial map $f : \Gamma^n \to \widehat{V}$.

We will later improve this to decomposition theorem over other bases: Every definable type on V can be decomposed into a definable type over Γ^n, and a germ of a definable function into $\widehat{V}$.

Exercise 3.9. Let M be a maximally complete model, and $\gamma \in \Gamma^n$. Then $M(\gamma) = \mathrm{dcl}(M \cup \{\gamma\})$ is algebraically closed.

Hint: Let N be a model containing $M(\gamma)$, and with $\Gamma(N)$ generated by γ over $\Gamma(M)$. For any $a \in N$, by Proposition 3.5, $tp(a/M(\gamma))$ extends to an $M(\gamma)$-definable type. In general if $e \in \mathrm{acl}(B)$ and $tp(e/B)$ extends to a B-definable type, show that $e \in \mathrm{dcl}(B)$.

Remark 3.10. Even over a base A consisting of imaginaries, if p is a stably dominated A-definable type and, then it is dominated by some A-definable function f into a finite-dimensional k-vector space. This follows from a general descent principle for stably dominated types and the elimination of imaginaries we will prove later.

3.11. **Definable modules.** We consider definable K-vector spaces $V \cong K^n$. When working over a base A we will always assume V has a basis of A-definable points; this can be taken as the definition, but in fact is automatic, at least over nontrivially valued subfields, by the following version of Hilbert 90:

Lemma 3.12. *Let F be a nontrivially valued field. If V is an F-definable K-space then V has a basis of F-definable points.*

Proof. We may assume $F = \mathrm{dcl}(F) \cap K$. In this case, F^{alg} is a model, so V has a basis of points of $V(F^{alg})$. This basis lies in $V(F')$ for some finite Galois extension F' of F. Now the automorphism group of F'/F in the sense of ACVF and of ACF coincide, by Lemma 3.13. Hence the usual Hilbert 90 applies. □

Lemma 3.13. *Let T be any expansion of the theory of fields, F a subfield of a model M of T with $\mathrm{dcl}(F) = F$. Let $F' \leq M$ be a finite normal extension of F. Then every field-theoretic automorphism of F'/F is elementary.*

Proof. Let G be the set of automorphisms of F'/F that are elementary, i.e. preserve all formulas. Then $Fix(G) = dcl(F) = F$. By Galois theory, $G = Aut(F'/F)$ in the field theoretic sense. □

Let Mod_V be the set of definable $\mathcal{O}$-submodules of V. $\Lambda \in Mod_V$ is *g-closed* if Λ intersects any 1-dimensional K-subspace $U \leq V$ in a submodule of the form $\mathcal{O}c$ or U or (0). Λ is a *semi-lattice* if it is g-closed and generates V as a K-space. Λ is a *lattice* if it is M-isomorphic to $\mathcal{O}^{\dim V}$.

Let V^* be the dual space to V; we identify V^{**} with V, and write (u, v) for the pairing $V \times V^* \to K$. For $\Lambda \in Mod_V$, let $\Lambda^* = \{v \in V^* : (\forall a \in \Lambda)(a, v) \in \mathcal{M}\}$. In class we considered a different notion, namely $\Lambda^*_c = \{v \in V^* : (\forall a \in \Lambda)(a, v) \in \mathcal{O}\}$.

Exercise 3.14. Let $\dim(V) = 1$, and $\Lambda \in Mod_V$. Then $\Lambda = \mathcal{O}c$ or $\Lambda = V$ or $\Lambda = (0)$ or $\Lambda = \mathcal{M}c$ for some $c \in M$.

Exercise 3.15. (1) * and $^*_c$ are weakly inclusion-reversing maps $Mod_V \to Mod(V^*)$. We have $\Lambda^{**} = \Lambda$, and if Λ is closed also $(\Lambda^*_c)^*_c = \Lambda$.
(2) Let $M \models ACVF$. If $\Lambda \in Mod_V(M)$ then Λ is M-isomorphic to $K^l \times \mathcal{O}^m \times \mathcal{M}^n$ for some l, m, n.
(3) If Λ is closed, then $\Lambda \cong K^l \times \mathcal{O}^m$ for some l, m.
(4) Λ^*_c is always closed.
(5) Define $\Lambda_c = (\Lambda^*_c)^*_c$. Then Λ is the smallest closed $\mathcal{O}$-module containing Λ. Λ contains $\mathcal{M}\Lambda$.

It follows from Example 3.15 (3) that the elements of Mod_V are uniformly definable.

Exercise 3.16. Let A be a valued field and let $e_1, \ldots, e_n$ be (imaginary) codes for modules, $\Gamma(A(e, \ldots, e_n) = \mathrm{dcl}(A \cup \{e_1, \ldots, e_n\}) \cap \Gamma$. Then there exists a maximally complete field N with $A \leq N$, $e_1, \ldots \in \mathrm{dcl}(A)$ and $\Gamma(N) = \Gamma(A(e, \ldots, e_n))$.

Hint: This reduces to the case $n = 1$, so $e = e_1$ codes a submodule Λ of K^n. We may assume Λ generates K^n, and the dual module generates the dual space; so Λ contains no nonzero subspace of K^n. Let Λ_c be the smallest lattice containing Λ. By adding to A a generic basis for Λ_c, we may assume $\Lambda_c = \mathcal{O}^n$. By Example 3.15 (5), $\mathcal{M}^n \subseteq \Lambda$. So to define Λ over A it suffices to define $\Lambda/\mathcal{M}^n$, a subspace of k^n. This can be done with parameters from k. If $\alpha \in K$, and a is a generic element of $\mathrm{res}^{-1}(\alpha)$, show that $\Gamma(A(a)) = \Gamma(A)$.

Exercise 3.17. Let Λ be a semi-lattice in V. For $a \in V$, show that $\{-\mathrm{val}(c) : ca \in \Lambda\}$ has a unique maximal element $v_\Lambda(a) \in \Gamma$, unless $a \subseteq \Lambda$; in the latter case write $v_\Lambda(a) = \infty$. Show that $v = v_\Lambda$ satisfies $v(a + b) \geq \min v(a), v(b)$ and $v(cb) = v(c) + v(b)$ for $a, b \in V, c \in K$. If Λ is a lattice, then (V, v_Λ) is a valued vector space. Conversely, given v with the above properties, $\Lambda_v = \{a : v(a) \geq 0\}$ is a semi-lattice, and alattice of $v(V \smallsetminus (0)) \subseteq \Gamma$.

3.18. **Pro-definable structure on $\widehat{V}$.** Let V be an affine variety, $V \subseteq \mathbb{A}^n$.

Let H_d be the space of polynomials in n variables of total degree $\leq d$. Let LH_d be set of semi-lattices in H_D.

Let $\widehat{V}$ denote the stably dominated types on V. We define $J_d : \widehat{V} \to LH_d$ by

$$J_d(p) = \{f \in H_d : (d_p v)(f(v) \in \mathcal{O})\}$$

$$J = (J_1, J_2, \ldots) : \widehat{V} \to \Pi_d LH_d$$

Proposition 3.19. (1) *J is 1-1.*
(2) *The image of J is a pro-definable set.*
(3) *In fact, the image of J_d is a definable set.*

(4) *Let $f(v,u)$ be a polynomial in variables $(v,u) = (v_1,\dots,v_n,u_1,\dots,u_m)$, of v-degree $\leq d$. There exists a definable function $h : LH_d \times \mathbb{A}^m \to \Gamma$ such that for any $p \in \widehat{V}$ and $b \in \mathbb{A}^m$, if $f_b(v) = f(v,b)$ then $\mathrm{val} f_b(v) = h(J_d(p),b)$ is in $p|\mathbb{U}$. In other words, $\mathrm{val} f_b(v)$ takes a constant value on generic realizations of p, this value as a function of p factors through $J_d(p)$, and it is uniformly definable over LH_d.*

Proof. It suffices to prove this for $V = \mathbb{A}^n$. Let $\bar{\Lambda} = (\Lambda_d)_{d\in\mathbb{N}} \in \Pi_d LH_d$.

Define $P(\bar{\Lambda}) = \{\mathrm{val}(f(x)) = v_{\Lambda_d}(x)\} : d \in \mathbb{N}, f \in H_d\}$, where v_{Λ_d} is as in Exercise 3.17.

Now check that $D = \{\bar{\Lambda} : P(\bar{\Lambda}$ is consistent $\}$ is a countable intersection of definable sets. If $P(\bar{\Lambda})$ is consistent, it generates a complete type over $\mathbb{U}$ (denoted the same way); type is always in $\widehat{V}$. Thus $J(\widehat{V}) = D$; this gives (2). Since $P(J(p))$ generates p, we have (1). With this definition of J, (4) is clear: $h(\Lambda,b) = v_\Lambda(f_b)$.

(3) is Theorem 3.1.1 in [3]; see a more explicit proof in the Appendix. □

4. Γ-internal subsets of $\widehat{V}$

Definition 4.1. A definable set D (possibly in imaginary sorts) is Γ-internal if (possibly over additional parameters) there exists a definable $Y \subset \Gamma^n$ and a surjective definable map $Y \to D$. Equivalently, there exists an injective definable map $D \to \Gamma^n$.

The equivalence in the definition uses elimination of imaginaries for Γ (an easy result.) In fact over one parameter from Γ, there even exist definable sets of representatives for any definable equivalence relation. Let $f : Y \to D$ be surjective. Let W be a definable set of representatives for the relation $f(y) = f(y')$. Then $g : D \to W$ defined by $f(g(d)) = d$ is a definable injective map.

We can call D *almost* Γ*-internal* if there exists a finite-to-one definable map $D \to \Gamma^n$. In fact by Example 3.9, almost Γ-internal definable sets are Γ-internal. For sets of lattices this can also be seen by noting that the proof of Proposition 4.6 goes through for almost Γ-internal sets, and that the conclusion implies Γ-internality.

If D is A-definable, it will turn out that the implicit parameters in the definition of Γ-internality can be taken to be in $\mathrm{acl}(A)$.

Lemma 4.2. *Let D be a Γ-internal subset of K^n. Then D is finite.*

Proof. It suffices to show that every projection of D to K is finite; so we may assume $n = 1$. If D is infinite, it contains an infinite closed ball; over additional parameters there is therefore a definable surjective map $D \to \mathrm{k}$. However if $Y \subseteq \Gamma^n$ there can be no surjective map $Y \to \mathrm{k}$, by the orthogonality of k, Γ. This contradiction shows that D is finite. □

Lemma 4.3. *Let D be a Γ-internal set of closed balls of equal radius in $\mathcal{O}$, i.e. $D \subseteq K/c\mathcal{O}$. Then D is finite.*

Proof. Let $D' = \cup D$. If D is infinite then D' contains a closed ball $d\mathcal{O} + e$ with $\mathrm{val}(d) < \mathrm{val}(c)$. Now $x \mapsto \mathrm{res}(d^{-1}(x-e))$ maps $d\mathcal{O}+e$ onto k, and factors through D. We obtain a contradiction as in Lemma 4.3. □

Lemma 4.4. *Let M be a model, $\gamma \in \Gamma^N$. Then any $M(\gamma)$-definable closed ball has a point in M.*

Proof. An $M(\gamma)$-definable closed ball b lies in some Γ-internal set D of closed balls. By Lemma 4.4, we may take D to be linearly ordered by inclusion. The intersection of all elements of D is a ball b' defined over M, closed or open, but nonempty; as M is a model, we can choose a point of b' over M. □

Lemma 4.5. *Any Γ-internal set D of balls is the union of a finite number of definable subsets, each linearly ordered by inclusion.*

Proof. Here we refer to Prop. 2.4.4 of [1]. □

We call a lattice Λ *diagonal* for a basis $(b_1, \ldots, b_n)$ if there exist $c_1, \ldots, c_n \in K$ with $\Lambda = \sum \mathcal{O} c_i b_i$. In other words, $\Lambda = \oplus_i \Lambda \cap K b_i$

Proposition 4.6. *Let M be a model. Let Λ be an $M(\gamma)$-definable lattice in $V = K^n$. Then Λ has an M-definable diagonalizing basis. Moreover if $e_1, \ldots, e_n$ is the standard basis, we can choose a diagonalizing basis of the form Ue, where strictly lower triangular matrix over M* [2]

Proof. The case $n = 1$ is trivial. Let V_1 be a one-dimensional subspace of V, $\mathbf{V} = V/V_1$, $g : V \to \mathbf{V}$ the canonical homomorphism. Choose b_1 such that $V_1 \cap \Lambda = \mathcal{O} b_1$. Let $\bar{\Lambda} = g\Lambda$. By induction, there exists an M-definable basis $\bar{b_2}, \ldots, \bar{b_n}$ diagonalizing $\bar{\Lambda}$; so $\bar{\Lambda} = \sum c_i \mathcal{O} \bar{b_i}$ for some c_i, with $c_i\mathcal{O}$ defined over $M(\gamma)$. Now $g^{-1}(c_i \bar{b_i}) \neq \emptyset$, and $g^{-1}(\bar{b_i})$ is a coset of V_1, so $c_i^{-1}\Lambda \cap g^{-1}(\bar{b_i})$ is a closed ball in V_1. By Lemma 4.4 it has an M-definable point b_i. Any element of Λ may be written as $v_1 + a_2 c_2 b_2 + \ldots + a_n c_n b_n$, with $v_1 \in V_1$, $a_i \in \mathcal{O}$. So $v_1 \in V_1 \cap \Lambda$. Thus $\Lambda = \oplus_i \Lambda \mathcal{O} b_i$. □

We may write $\Lambda = \oplus_{i=1}^n \mathcal{O}\gamma_i U e_i = U \oplus_{i=1}^n \mathcal{O}\gamma_i e_i = U S_\gamma \mathcal{O}^n$, where S_γ is the diagonal matrix $(\gamma_1, \ldots, \gamma_n)$.

Given Λ, the matrix US_γ is determined up to multiplication on the right by an element of $B_n(\mathcal{O})$; and S_γ is determined by US_γ; the image of S_γ in $D_n/D_n(\mathcal{O}) = \Gamma^n$ depends on Λ alone, and we denote it $G(\Lambda)$. (This corresponds to the composed homomorphism $\gamma : B_n \to \Gamma^n$, composition of $B_n \to B_n/U_n = D_n$ with the natural map $D_n \to \Gamma^n$.)

[2] 'strict' here means: 1's on the diagonal.

Corollary 4.7. *let D be a Γ-internal set of lattices. Then there exist a finite partition $D = \cup_{i=1}^{r} D_i$ and bases $b^1, \ldots, b^r$ such that each $\Lambda \in D_i$ is diagonal in b^i. The bases b^i are strictly upper triangular. The function G defined in the paragraph above is injective on each D_i.*

Proof. As the matrix U in the conclusion of Proposition 4.6 is defined over M, while Λ varies over a definable set, compactness assures the existence of finitely many matrices $U_1, \ldots, U_r$ over M, such that each $\Lambda \in D$ has the form $U_i S_\gamma \mathcal{O}^n$ for some $i \leq r$ and for $\gamma = G(\Lambda)$. Let $D_i = U_i D_n \mathcal{O}^n$.

□

5. Definable types in ACVF

Let M be a model. We say that $tp(a/M)$ is definable if there exists a (necessarily unique) M-definable type p with $tp(a/M) = p|M$.

Lemma 5.1. *If $tp(a/M)$ is definable, and $c \in \mathrm{acl}(Ma)$, then $tp(ac/M)$ is definable.*

Proof. Let $\phi(xy) \in tp(ac/M)$ be a formula such that $\phi(a, y)$ has m solutions, with m least possible. Then $p(x)|N \cup \phi(x, y)$ generates a complete type over any elementary extension N. By Lemma 1.9, this is a definable type. □

Lemma 5.2. *Let A be any subset of $\mathbb{U}^{eq}$, i.e. any set consisting possibly of imaginary elements. Let $V \subseteq K^n$ be an A-definable set. Then there exists a definable type on V, over $\mathbb{U}$, with finite orbit under $Aut(\mathbb{U}/A)$.*

This comes as close as possible to saying that p is A-definable; one cannot do better since V might be finite, or may have a finite but nontrivial definable quotient.

Proof. By induction on n. If $n = 1$, V contains finitely many balls, each with some finite union of sub-balls missing. The generic type of one of these balls will do. For $n > 1$, let $\pi : K^n \to K^{n-1}$ be the projection, and let p' be a definable type on $V' = \pi(V)$ with finite orbit. Let M be a model containing A, and let $a \models p'|M$. Let p'' be a definable type on $\pi^{-1}(a)$ with finite orbit under $Aut(\mathbb{U}/A(a))$. So p'' is $A(a')$ definable, with $a' \in acl(A(a))$. Let $a'' \models p''|M(a, a')$. By Lemma 5.1, $tp(aa'/M)$ is definable, and hence $tp(aa'a''/M)$ is definable, so $tp(aa''/M)$ is definable, i.e. equals $p|M$ for some definable type p. The number of conjugates of p is at most the number of conjugates of $a'/A(a)$. □

Let r be an A-definable type on Γ^n. By a *pro-definable function on r into $\widehat{V}$* we mean a pro-definable function f represented by a sequence of definable functions f_i, such that $\mathrm{dom}(f_i) \in r|A$ for each i.

Let f be an pro-definable function on r into $\widehat{V}$ with $\mathrm{dom}(f) \in r|A$, whose p-germ is defined over A. Recall the definition of $\int_r f$ (Example 3.7). It depends on f only through the p-germ of f, so that $\int_r f$ is an A-definable type.

Theorem 5.3. *Let p be an A-definable type on a variety V. Then there exist a definable type r on Γ^n and a definable r-germ f of pro-definable maps into $\widehat{V}$, with $p = \int_r f$.*

Proof. Let M be a maximally complete model, containing A.

Let $c \models p|M$. Let $\gamma = g'(c)$ be a basis for $\Gamma(M(c))$ over $\Gamma(M)$; let $r' = g'_* p$.

Now $tp(c/M(\gamma))$ is stably dominated, so it equals $q|M(\gamma)$ for some $q \in \widehat{V}$; we can write $q = f'(\gamma)$, with f' an M-definable function into $\widehat{V}$. By definition, $p = \int_{r'} f'$.

□

The proof showed that $r = g_* p$, where $g = \alpha^c \circ g'$. In particular, the r-germ of $g \circ f$ is the r-germ of the identity, i.e. f is genericallly injective. (We could also arrange this a posteriori.)

How canonical is the pair (r, f)?

Definition 5.4. Consider pairs (r, h) with r a definable type and h a definable function. We say two such pairs $(r, h), (r', h')$ are *equivalent up to generic reparameterization*, $(r, h) \sim (r', h')$, iff there exist definable functions ϕ, ϕ' such that $\phi_* r = \phi'_* r'$, and for some definable h'', $h = h'' \circ \phi$ and $h' = h'' \circ \phi'$.

When h' is generically injective, this is equivalent to the existence of a a definable ϕ such that $r' = \phi_* r$ and $h = h' \circ \phi$ as an r-germ.

If h is pro-definable, with target $X = \varprojlim X_k$ and $\pi_k : X \to X_k$ the defining maps, we say $(r, h) \sim (r', h')$ if $(r, \pi_k \circ h) \sim (r', \pi_k \circ h')$ for each k.

Lemma 5.5. *The pair (r, f) is determined by $p = \int_r f$, up to generic reparameterization.*

Proof. Suppose $p = \int_r f = \int_{r'} f'$, with r, r', f', f' defined over some N. Let $\gamma models r|N$, $c \models f(\gamma)|N(\gamma)$. So $c \models p|N$. Since also $p = \int_{r'} f'$, we may find $\gamma' \models r'|N$ such that $c \models f'(\gamma')|N(\gamma')$. By stable domination of p, we have $\Gamma(N(c)) \subset N(\gamma)$. We claim that $\gamma \in \Gamma(N(c))$. Let γ'' be a basis for $\Gamma(N(c))$ over N. Then $tp(c/N(\gamma''))$ extends to a stably dominated type p'' defined over $N(\gamma'')$. By orthogonality to Γ again, p'' implies a complete type over $N(\gamma)$, namely $tp(c/N(\gamma)) = p$. It follows that $p = p''$ is based on $N(\gamma'')$, and so by generic injectivity of f we have $\gamma' \in \mathrm{dcl}(N(\gamma''))$. Thus $N(\Gamma(N(c))) = N(\gamma)$ and similarly $N(\Gamma(N(c))) = N(\gamma')$. So $N(\gamma) = N(\gamma')$. Moreover $h(\gamma), h'(\gamma')$ are stably dominated types based on $N(\gamma)$ and with the same restriction to this base, namely $tp(c/N(\gamma))$; so $h(\gamma) = h'(\gamma')$. Let ϕ be an invertible N-definable function such that $\gamma' = \phi(\gamma)$; then $r' = \phi_* r$ and as $h'(\gamma') = h(\phi^{-1}(\gamma'))$, $h' = h \circ \phi^{-1}$. □

We will study this notion in the ACVF setting in the next section, but we indicate now how it will go. We will see in Lemma 6.2 that after a possible reparametrization, one can find an A-definable function G on $\widehat{V}$ such that $G \circ$

f is the identity germ on r. (Basically this is the 0-definable function G of Corollary 4.7; we need A only in order to find an affine patch V' of V and identify $\widehat{V'}$ with a sequence of lattices.) This implies that $r = G_*p$ is A-definable, and also rigidifies f so that reparameterization is no longer possible, and the r-germ of f is uniquely determined. *Hence with these choices we find an A-definable r and a function f with A-pro-definable germ.* We can even use Lemma 1.21 to make r, if we wish, 0-definable; this requires an additional reparamterization by a certain A-definable translation.

Remark 5.6. Though the r-germ of f can be chosen to be A-pro-definable, it is not always possible to find an A-(pro)definable f. For instance for the generic type of an A-definable open ball without an A-definable sub-ball, this is the case. This phenomenon is responsible for much of the subtlety in the stability-theoretic study of ACVF.

The function G described above, inverting the germ f on the left, cannot in general be take of the form $p \mapsto g_*p$ for any A-definable g.

6. Imaginaries in ACVF

Recall B_n denote the group of invertible upper triangular matrices. U_n is the group of matrices in B_n with 1's on the diagonal. D_n is the group of diagonal matrices, so that $B_n = D_nU_n$.

If G is any algebraic subgroup of the group GL_n of invertible $n \times n$ - matrices, $G(\mathcal{O})$ denotes the elements $M \in G$ such that M, M^{-1} have entries in $\mathcal{O}$.

Let S_n be the coset space $B_n/B_n(\mathcal{O})$. We will see below that any lattice in K^n has a triangular basis. Hence B_n acts transitively on the set of lattices; and $B_n(\mathcal{O})$ is the stabilizer of the standard lattice $\mathcal{O}^n$. It follows that $B_n/B_n(\mathcal{O})$ can be identified with the set of lattices in K^n. (By a similar argument, so can $GL_n(K)/GL_n(\mathcal{O})$.)

Let $\widetilde{GL_n}(\mathcal{O})$ be the pullback of the stabilizer of a vector, under the natural homomorphism $GL_n(\mathcal{O}) \to GL_n(\mathrm{k})$. Let T_n be the coset space $GL_n/\widetilde{GL_n}(\mathcal{O})$ We have a natural map $T_n \to S_n$. Given $b \in S_n$, viewed as a lattice Λ, naming an element of T_n is equivalent to choosing a point of $\Lambda/\mathcal{M}\Lambda$. Let GG consist of the valued field sort K, along with the sorts S_n, T_n.

Certain related imaginary sorts can be directly shown to be coded in the sorts S_n, T_n.

Lemma 6.1. (1) *Any definable $\mathcal{O}$-submodule of K^n, as well any coset of such a submodule of K^n, can be coded in GG.*
(2) *Any finite subset of $S_n \cup T_n \cup K^m$ is coded in GG.*
(3) *Let H be a subgroup of U_n defined by a conjunction*

$$H = \{a \in U_n : \bigwedge_{i \le j \le n} \mathrm{val}(a_{ij}) \diamond_{ij} \alpha_{ij}\}$$

where $\alpha_{ij} \in \Gamma_\infty$ and $\diamond$ denotes $\geq$ or $>$. Let A be a base structure containing $\alpha_{ij}, i, j \leq n$. Then any coset of H is coded in GG_A (i.e. for any coset C of H there exists $g \in GG^m$ such that g is a canonical code for C over A.)

Proof. We will not repeat the proofs of (1,2) from [1]; (1) is rather straightforward, see 2.6.6; (2) is Prop. 3.4.1 there. .

For (3), let A_n be the $\mathcal{O}$-algebra of strict [3] upper triangular matrices. Let J be the subalgebra defined by: $\bigwedge_{i \leq j \leq n} \mathrm{val}(a_{ij}) \diamond_{ij} \alpha_{ij}$. Then $H = 1 + J$. We have $aH = bH$ iff $a = b(1 + j)$ for some $j \in J$ iff $aJ = bJ =: J'$ and $a + J' = b + J'$. As J' is an $\mathcal{O}$-module and $a + J'$ a coset, (3) follows from (1). □

Lemma 6.2. *Let r be a definable type on a definable $D \subset \Gamma^n$, $V = K^N$, and $h : D \to LV$ be an injective definable map. Then $(r, h)/\sim$ has a canonical base in GG.*

Proof. Let $U(\Lambda)$ be the maximal K-subspace contained in $\Lambda \in L$. Say $\dim U(h(t)) = d$. $U(h(t))$ can be viewed as an element of a Grassmanian variety $Gr_d(V)$. By Lemma 4.2, the image of $U(h(t))$ is finite. Since p is complete, the image is a single element U, i.e. $U(h(t)) = U_d$ for all $t \models r$. Now U is clearly an invariant of $(r, h)/\sim$. We may work over a base where all U are defined, and view h as a function $r \to L(V/U)$. We may thus assume $h(t)$ is a lattice for $t \models r$.

By Corollary 4.7 there exists a triangular basis b for V such that $h(t)$ is diagonal in $b = (b_1, \ldots, b_n)$, for $t \models r$. So $h(t) = \sum \mathcal{O}\gamma_i(t) b_i$ for certain definable functions $\gamma_i : r \to \Gamma$. Let $\gamma = (\gamma_1, \ldots, \gamma_n)$. (Recall $\mathcal{O}\gamma$ denotes $\{x : \mathrm{val}(x) \geq \gamma\}$.) We can replace r by $\gamma_* r$ and h by the function $h_b((s_1, \ldots, s_n)) = \sum \mathcal{O} s_i b_i$, without changing the $\sim$-class. So from now on we will consider only $h = h_b$ of this form. Thus we need to code pairs (r, b) up to $\sim$, where $(r, b) \sim (r', b')$ iff $(r, h_b) \sim (r', h_{b'})$. Note that h_b is injective on Γ^n.

By Lemma 1.21 there exists $c \in \Gamma^n$ such that $\alpha^c_*(r)$ is 0-definable; where α^c is translation by c. Say $c = (c_1, \ldots, c_n), c_i = -\mathrm{val}(e_i)$. Let e be the diagonal matrix $(e_1, \ldots, e_n)$. Then $(r, b) \sim (\alpha^c_*(r), eb)$. Replacing (r, b) by $(\alpha^c_*(r), eb)$, we may assume r is 0-definable.

Since r is 0-definable, $(r, b)/\sim$ is equi-definable with $b/\sim$, so we will now fix r and consider the equivalence relation: $b \sim b'$ iff $(r, b) \sim (r, b')$

We view $b = (b_1, \ldots, b_n)$ as a matrix, with b_i the i'th column. Note that b, b' generate the same $\mathcal{O}$-lattice iff $b'_i = \sum c_{ij} b_j$ for some $c_{ij} \in \mathcal{O}$ and conversely, iff $b' = bN$ for some $N \in GL_n(\mathcal{O})$. Also, we have $(t_1 b_1, \ldots, t_n b_n), bD_t$ generate the same $\mathcal{O}$-module, where D_t denotes any triangular matrix $(e_1, \ldots, e_n)$ with $\mathrm{val}(e_i) = t_i$.

[3] 'strict' here means: 0's on the diagonal

Suppose $b \sim b'$. So $r' = \phi_* r$ for some definable function ϕ and $h_b = h_{b'} \circ \phi$. Let $N = N(b, b')$ be the change of basis matrix, $bN(b, b') = b'$. Then $N(b, b')$ is upper triangular. Write $s = \phi(t)$. Then bD_t and $b'D_s = bND_s$ generate the same $\mathcal{O}$-module, so $bD_t = bND_sN'$ for some $N' \in GL_n(\mathcal{O})$, or $D_s^{-1}ND_t \in GL_n(\mathcal{O})$. Equivalently $D_s^{-1}ND_t \in B_n(\mathcal{O})$. But N is upper triangular, $N(b, b') = D(b, b')U(b, b')$, with $D(b, b')$ diagonal and $U(b, b')$ strictly upper triangular. It follows that $D(b, b') = D_{t-s} \mod D(\mathcal{O})$. This holds for $t \models r$; so $t - s$ is constant, i.e. $\phi(t) = t + c_0$, where $c_0 \in S(r) = \{c \in \Gamma^n : \alpha_*^c r = r\}$. Note that $S(r)$ is a definable subgroup of Γ^n (of the form $E(\Gamma^l \times (0))$ for some $l \leq n$ and some matrix E with $\mathbb{Q}$-coefficients.) Let $S'(r)$ be the pullback of $S(r)$ to the group $D_n(K)$ of diagonal matrices. Then $D(b, b') \in S'(r)$. Moreover, since $D_s^{-1}ND_t \in B_n(\mathcal{O})$, we have $U(b, b') \in D_t B_n(\mathcal{O}) D_t^{-1}$, or $U(b, b') \in D_t U_n(\mathcal{O}) D_t^{-1}$.. Conversely, the argument reverses to show that if $D(b, b') \in S'(r)$ and $U(b, b') \in D_t U_n(\mathcal{O}) D_t^{-1}$ for generic $t \models r$, then $b \sim b'$. Let $D_\nu = \{g \in U_n : (d_r t)(g \in D_t U_n(\mathcal{O}) D_t^{-1})\}$. It is easy to see that this is one of the groups in Lemma 6.1 (3), and hence coded in GG.

□

Theorem 6.3. *In the sorts GG, ACVF admits elimination of imaginaries.*

Proof. By Lemma 5.2, Lemma 1.17 and Lemma 6.1 (2), it suffices to show that any definable type q on $V = \mathbb{A}^n$ has a canonical base in the sorts GG. Now q has the form $\int_r h$ where r is a definable type on Γ^m and $h : \Gamma^m \to \widehat{V}$ is a definable map. q is equi-definable with the pair (r, h) up to generic reparameterization.

We have $h = (h_d)$, $h_d : r \to LH_d$, where H_d is the space of polynomials in n variables of degree $\leq d$. Define $\sim_d$ as $\sim$ above. For large enough d, h_d is injective on a definable neighborhood of r. [4] Clearly if σ fixes q then it fixes $(r, h_d)/\sim_d$ for each d; conversely if σ fixes $(r, h_d)/\sim_d$ for large enough d, then it fixes the q-definition of any given formula, so it fixes q. Thus it suffices to code $(r, h_d)/\sim_d$ for each d. This was proved in Lemma 6.2. □

7. Appendix

We give here an effective description of the image of $\widehat{V}$ in the space of semi-lattices. This description came out of a conversation with Bernd Sturmfels.

Let F be a valued field. We will use Robinson's quantifier-elimination theorem in a two-sorted version, i.e. some variables range over K and some range over the residue field k. This follows easily from the one-sorted version: if $\phi(x_1, \ldots, x_n)$ is quantifier-free formula on $\mathcal{O}^n$, which is invariant under translation by $\mathcal{M}^n$, then the solution set of ϕ can be viewed as a subset of k^n; and it is easy to see that this subset is constructible (a Boolean combination of varieties.) Note that if

[4] alternatively, for any d, we can factor out the kernel of h_d and work with the pushforward r_d of r.

$\phi(x_1, \ldots, x_n, y_1, \ldots, y_m)$ is $\mathcal{M}^n$-translation invariant in the x-variables, with the y-variables fixed, then so is $(\exists y_1)(\exists y_2)\phi$, or any other sequence of quantifiers over the y-variables.

7.1. Let us take an affine variety $V = Spec(k[X_1, \ldots, X_m]/I)$. Let $\widehat{V}$ be the stable completion. Let H_d be the vector space of polynomials of degree $\leq d$, $I_d = H_d \cap I$, $U_d = H_d/I_d$, and let $S(U_d)$ be the space of semi-lattices in U_d.

There is a natural map $r_{d,n} : \widehat{V} \to S(U_d)$. Namely if p is viewed as a semi-norm, $r_{d,n}(p) = \{f + I_d : p(f) \geq 0\}$.

Given $\Lambda \in S(U_d)$, let $T = T(\Lambda)$ be the maximal K-space contained in Λ, and let $(f_1, \ldots, f_n)$ be an $\mathcal{O}$-basis for Λ/T. Let

$$R(\Lambda) = \{(\mathrm{res}f_1(c), \ldots, \mathrm{res}f_n(c)) : c \in V(K), v(f(c)) \geq 0 \text{ for } f \in \Lambda\}$$

By Robinson's theorem, this is a constructible subset of k^n. If we change the $\mathcal{O}$-basis, $R(\Lambda)$ changes by a linear transformation.

Lemma 7.2. *$\Lambda \in r_{d,n}(\widehat{V})$ iff $R(\Lambda)$ is not contained in a finite union of proper subspaces of k^n.*

Proof. First suppose $\Lambda \in r_{d,n}(\widehat{V})$; say $\Lambda = r_{d,n}(p)$. Suppose $R(\Lambda)$ is contained in a finite union of proper subspaces of k^n; these subspaces and all data are defined over some model M. Let $c \models p|M$; let $f_1, \ldots, f_n$ be an $\mathcal{O}$-basis for Λ; then $(\mathrm{res}f_1(c), \ldots, \mathrm{res}f_n(c)) \in R(\Lambda)$, so it must lie in one of the M-definable proper subspaces mentioned above; i.e. $\sum \alpha_i \mathrm{res}f_i(c) = 0$, $\alpha_i \in k(M)$, not all 0. Extend $(\alpha_1, \ldots, \alpha_n)$, viewed as an element χ_1 of $(k^n)^*$, to a basis $\chi_1, \ldots, \chi_n$ of $(k^n)^*$, and lift to a basis $\chi_1, \ldots, \chi_n$ of dual lattice Λ^* of Λ. Let $g_1, \ldots, g_n$ be the dual basis of Λ. Then whenever $a \models p$, $g_1(a)$ has positive valuation; say $\alpha = \mathrm{val}(c)$; it follows that $c^{-1}g_1 \in \Lambda$, but $c^{-1} \notin \mathcal{O}$, a contradiction.

Conversely, assume $R(\Lambda)$ is not contained in a finite union of proper subspaces of k^n. Let M be a maximally complete model over which V, Λ are defined, let $f = (f_1, \ldots, f_n)$ be a basis for Λ over M. Find $c \in V$ be such that $f(c) = 0$ for $f \in T$, $\mathrm{val}f_i(c) \geq 0$ for $i \leq n$ and $(\mathrm{res}f_1(c), \ldots, \mathrm{res}f_n(c))$ does not lie in any proper M-definable subspace of k^n. Let α be a basis for $\Gamma(M(c))$ over $\Gamma(M)$; so there exists a stably dominated type p over $M(\alpha)$ with $c \models p|M(\alpha)$. It is clear that $T \subseteq I(p) \cap H_d$ (where $I(p)$ is the kernel of the semi-valuation p) and $\Lambda \subset r_{d,n}(p)$. We claim that in fact, $r_{d,n}(p) = \Lambda$. For suppose (e.g.) that $I_d = I(p) \cap H_d$ but $r_{d,n}(p)$ is a bigger lattice Λ'. The lattice Λ' is defined over $M(\alpha)$ and so lies in a Γ-parameterized family of lattices over M, so there exists a basis $g_1, \ldots, g_n$ of H_d/T such that Λ' is diagonal in this basis, i.e. Q is generated by $c_1g_1, \ldots, c_ng_n$ for some $c_1, \ldots, c_n$. The change-of-basis matrix Q from $c_1g_1, \ldots, c_ng_n$ to $f_1, \ldots, f_n$ lies in $M_n(\mathcal{O})$; if it is in $GL_n(\mathcal{O})$, then the lattices are equal; if not, then some element e of $\Lambda(M)$ lies in $\mathcal{M}\Lambda'$ but not in $\mathcal{M}\Lambda$. As $e \notin \mathcal{M}\Lambda$, we have $\mathrm{res}e(c) \neq 0$,

otherwise $(\mathrm{res}f_1(c_m), \ldots, \mathrm{res}f_n(c_m))$ would lie in a proper subspace. It follows that $\mathrm{val}e(c) = 0$, and so $p(e) = 0$, contradicting $e \in \mathcal{M}\Lambda'$.

□

Now one can algorithmically decompose the constructible set $R(\Lambda)$ into irreducible, relatively closed sets and find their linear span; the condition of the lemma is that one of these spans should have dimension n. This gives an effective description of the image of $\widehat{V}$.

7.3. We have in general $\dim(R(\Lambda)) \leq \dim(V)$. An important subset of the stable completion (denoted $V^{\#}$) consists of the strongly stably dominated points (see [3]). In the present setting, a stably dominated type on a variety V is strongly stably dominated iff the residue field extension it induces has the same transcendence degree as the field extension it induces.

Now if Λ is a lattice with $\dim(R(\Lambda)) = \dim(V)$, then Λ is the image of at most a finite number $n(\Lambda)$ of elements p of $\widehat{V}$, such that for $g_1, \ldots, g_n$ a basis of Λ, M a model over which the data is defined, and $c \models p|M$, $\mathrm{res}g_1(c), \ldots, \mathrm{res}g_n(c)$ are linearly independent over $\mathrm{k}(M)$. These points p all lie in $V^{\#}$; and an upper bound on their number is easily given. This raises the hope of describing elements of $V^{\#}$ via a single tropical approximation. But we have:

Problem 7.4. Let Λ be given, and assume $\dim(R(\Lambda)) = \dim(V)$. Determine $n(\Lambda)$ (or just whether $n(\Lambda) = 1$) effectively.

References

[1] D. Haskell, E. Hrushovski, D. Macpherson, *Definable sets in algebraically closed valued fields: elimination of imaginaries*, J. Reine Angew. Math. **597** (2006), 175–236.

[2] D. Haskell, E. Hrushovski, D. Macpherson, *Stable domination and independence in algebraically closed valued fields*, Lecture Notes in Logic, 30. Association for Symbolic Logic, Chicago, IL; Cambridge University Press, Cambridge, 2008.

[3] E. Hrushovski., F. Loeser, Non-Archimedean tame topology and stably dominated types, arXiv:1009.0252

[4] E. Hrushovski, Definable sets over valued fields, Valuation Theory conference, El Escorial, July 2011, `http://math.huji.ac.il/ ehud/escorial.pdf`

[5] Kamensky, Moshe Ind- and pro- definable sets. Ann. Pure Appl. Logic 147 (2007), no. 3, 180186.

[6] A. Pillay, *Model theory and stability theory, with applications in differential algebra and algebraic geometry*, in Model theory and Applications to Algebra and Analysis, volume 1, LMS Lecture Notes Series 349, 2008 (edited by Chatzidakis, Macpherson, Pillay, Wilkie), 1– 23. See also Lecture notes on Model Theory, Stability Theory, Applied Stability theory, on http://www.maths.leeds.ac.uk/∼pillay.

[7] L. van den Dries, *Dimension of definable sets, algebraic boundedness and henselian fields*, Ann. Pure Appl. Logic **45** (1989), 189–209.

[8] L. van den Dries, *Tame topology and o-minimal structures*, Cambridge Univ. Press, New York, 1998.

Institute of Mathematics, the Hebrew University of Jerusalem, Givat Ram, Jerusalem, 91904, Israel.

E-mail address: ehud@math.huji.ac.il